

CLAIMS

- 1 1. A network interface device, comprising:
2 receive logic, which is coupled to receive from a
3 network a sequence of data packets, each packet
4 comprising respective header data;
5 a protocol processor, coupled to read and process
6 the header data so as to identify a group of the received
7 packets that contain respective fragments of a data
8 frame, the fragments having a fragment order within the
9 data frame; and
10 host interface logic, which is coupled to a host
11 memory accessible by a host processor, and is controlled
12 by the protocol processor so as to allocate space for the
13 data frame in the host memory, and to reassemble the
14 fragments of the data frame in the fragment order in the
15 space allocated in the host memory.
- 1 2. A device according to claim 1, wherein the protocol
2 processor is arranged to identify the group of the
3 packets and to control the host interface logic so as to
4 reassemble the fragments irrespective of whether the
5 sequence in which the packets are received coincides with
6 the fragment order.
- 1 3. A device according to claim 2, wherein the protocol
2 processor is arranged to control the host interface logic
3 so as to allocate the space for the data frame responsive
4 to the header data of the packet received first in the
5 sequence among the packets in the group, regardless of
6 whether the packet received first in the sequence is
7 first in the fragment order.

- 1 4. A device according to claim 1, wherein the protocol
2 processor is arranged to control the host interface logic
3 so as to write each of the fragments to a respective
4 location within the allocated space in the host memory
5 responsive to a fragment offset parameter in the header
6 data of each of the packets.
- 1 5. A device according to claim 1, wherein the protocol
2 processor is arranged to control the host interface logic
3 so as to allocate the space for the data frame responsive
4 to the packet that is received first in the sequence
5 among the packets in the group.
- 1 6. A device according to claim 5, wherein the protocol
2 processor is arranged to determine a size of the space to
3 allocate responsive to a frame length indication in the
4 header data of the packet received first in the sequence.
- 1 7. A device according to claim 6, wherein the frame
2 length indication comprises one or more fields in the
3 header data indicating an exact length of the data frame.
- 1 8. A device according to claim 6, wherein the frame
2 length indication comprises one or more fields in the
3 header data indicating an upper bound on a length of the
4 data frame.
- 1 9. A device according to claim 1, wherein the protocol
2 processor is arranged to monitor a time required to
3 receive all of the packets in the group, and to control
4 the host interface logic so as to release the space
5 allocated for the data frame if the time exceeds a
6 predetermined limit without all of the fragments in the
7 data frame having been reassembled.

1 10. A device according to claim 9, wherein when the time
2 exceeds the predetermined limit without all of the
3 fragments in the data frame having been reassembled, the
4 protocol processor is arranged to return a message over
5 the network to a source of the packets indicating that
6 the data frame was not received.

1 11. A device according to claim 1, wherein the group of
2 the received packets is one of a plurality of different
3 groups, the packets in the different groups containing
4 the fragments of different, respective data frames, and
5 wherein the protocol processor is arranged to identify
6 the different groups and to control the host interface
7 logic so as to simultaneously reassemble the fragments of
8 the different data frames in respectively-allocated
9 spaces in the host memory.

1 12. A device according to claim 1, wherein the packets
2 are transmitted over the network and received by the
3 receive logic in accordance with one or more
4 communication protocols, whereby the header data in the
5 packets comprise protocol information, and wherein the
6 protocol processor is arranged to identify the group and
7 to control the host interface logic responsive to the
8 protocol information.

1 13. A device according to claim 12, wherein the one or
2 more communication protocols comprise a plurality of
3 different communication protocols, and wherein the
4 protocol processor is arranged to select the group of
5 packets for reassembly depending on which of the
6 communication protocols was used in transmitting the
7 packets.

1 14. A device according to claim 12, wherein the one or
2 more protocols comprise a network layer protocol.

1 15. A device according to claim 14, wherein the network
2 layer protocol comprises an Internet Protocol (IP).

1 16. A device according to claim 12, wherein the one or
2 more protocols comprise a transport layer protocol.

1 17. A device according to claim 16, wherein the
2 transport layer protocol comprises a Transport Control
3 Protocol (TCP).

1 18. A device according to claim 16, wherein the
2 transport layer protocol comprises a User Datagram
3 Protocol (UDP).

1 19. A device according to claim 1, wherein the receive
2 logic is coupled to receive the packets over the network
3 from a plurality of different sources, and wherein the
4 protocol processor is arranged to select the group of
5 packets for reassembly dependent on the packets having
6 been received from one or more chosen sources among the
7 plurality of different sources.

1 20. A device according to claim 19, wherein the selected
2 sources are chosen responsive to a level of reliability
3 of a connection over the network between the chosen
4 sources and the network interface device.

1 21. A device according to claim 1, wherein the protocol
2 processor is arranged to control the host interface logic
3 so as to write the data packets that do not belong to the
4 identified group to the host memory substantially without
5 reassembly processing thereof by the network interface
6 device.

1 22. A device according to claim 1, wherein the host
2 interface logic comprises a direct memory access (DMA) ~
3 engine, which is arranged to write the fragments to the
4 host memory substantially without involvement of the host
5 processor.

1 23. A device according to claim 22, wherein the host
2 interface logic is coupled to notify the host processor
3 that the fragments of the data frame have been
4 reassembled in the host memory only after all of the
5 fragments have been reassembled.

1 24. A device according to claim 23, wherein the protocol
2 processor is arranged to determine a total length of the
3 data frame responsive to the header data of at least one
4 of the packets in the group, and to count a quantity of
5 the data in the fragments reassembled in the data frame,
6 and to determine that all of the fragments have been
7 reassembled by comparing the total length to the quantity
8 of the data reassembled.

1 25. A device according to claim 23, wherein the protocol
2 processor is further arranged to control the host
3 interface logic to write a frame header to the allocated
4 space in the host memory, indicating to the host
5 processor a total length of the data frame.

1 26. A device according to claim 1, wherein the receive
2 logic, protocol processor and host interface logic are
3 contained together in a single integrated circuit chip,
4 which is separate from the host processor and host
5 memory.

1 27. A method for interfacing a host processor to a
2 network, comprising:

3 receiving a sequence of data packets from the
4 network at a network adapter, each packet comprising
5 respective header data;

6 processing the header data in the adapter so as to
7 identify a group of the received packets that contain
8 respective fragments of a data frame, the fragments
9 having a fragment order within the data frame;

10 allocating space for the data frame in a host memory
11 accessible by the host processor;

12 writing the fragments of the data frame from the
13 network adapter to the space allocated in the host memory
14 so that the fragments are reassembled in the space in the
15 fragment order; and

16 notifying the host processor when all of the
17 fragments have been reassembled.

1 28. A method according to claim 27, wherein processing
2 the header data comprises identifying the group of the
3 packets and preparing the fragments for writing to the
4 host memory irrespective of whether the sequence in which
5 the packets are received coincides with the fragment
6 order.

1 29. A method according to claim 28, wherein allocating
2 the space comprises allocating the space for the data
3 frame responsive to the header data of the packet
4 received first in the sequence among the packets in the
5 group, regardless of whether the packet received first in
6 the sequence is first in the fragment order.

1 30. A method according to claim 27, wherein writing the
2 fragments comprises writing each of the fragments to a
3 respective location within the allocated space in the

4 host memory responsive to a fragment offset parameter in
5 the header data of each of the packets.

1 31. A method according to claim 27, wherein allocating
2 the space comprises allocating the space for the data
3 frame responsive to the packet that is received first in
4 the sequence among the packets in the group.

1 32. A method according to claim 31, wherein allocating
2 the space comprises determining a size of the space to
3 allocate responsive to a frame length indication in the
4 header data of the packet received first in the sequence.

1 33. A method according to claim 32, wherein the frame
2 length indication comprises one or more fields in the
3 header data indicating an exact length of the data frame.

1 34. A method according to claim 32, wherein the frame
2 length indication comprises one or more fields in the
3 header data indicating an upper bound on a length of the
4 data frame.

1 35. A method according to claim 27, and comprising
2 monitoring a time required to receive all of the packets
3 in the group, and releasing the space allocated for the
4 data frame if the time exceeds a predetermined limit
5 without all of the fragments in the data frame having
6 been reassembled.

1 36. A method according to claim 35, and comprising, when
2 the time exceeds the predetermined limit without all of
3 the fragments in the data frame having been reassembled,
4 returning a message over the network to a source of the
5 packets indicating that the data frame was not received.

1 37. A method according to claim 27, wherein the group of
2 the received packets is one of a plurality of different

3 groups, the packets in the different groups containing
 4 the fragments of different, respective data frames, and
 5 wherein processing the header data comprises identifying
 6 the different groups so as to simultaneously reassemble
 7 the fragments of the different data frames in
 8 respectively-allocated spaces in the host memory.

1 38. A method according to claim 27, wherein the packets
 2 are transmitted over the network and received by the
 3 receive logic in accordance with one or more
 4 communication protocols, whereby the header data in the
 5 packets comprise protocol information, and wherein
 6 processing the header data comprises identifying the
 7 group and preparing the fragments for writing responsive
 8 to the protocol information.

1 39. A method according to claim 38, wherein the one or
 2 more communication protocols comprise a plurality of
 3 different communication protocols, and wherein
 4 identifying the group comprises selecting the group of
 5 packets for reassembly depending on which of the
 6 communication protocols was used in transmitting the
 7 packets.

1 40. A method according to claim 38, wherein the one or
 2 more protocols comprise a network layer protocol.

1 41. A method according to claim 40, wherein the network
 2 layer protocol comprises an Internet Protocol (IP).

1 42. A method according to claim 38, wherein the one or
 2 more protocols comprise a transport layer protocol.

1 43. A method according to claim 42, wherein the
 2 transport layer protocol comprises a Transport Control
 3 Protocol (TCP).

1 44. A method according to claim 42, wherein the
2 transport layer protocol comprises a User Datagram -
3 Protocol (UDP).

1 45. A method according to claim 27, wherein receiving
2 the sequence of data packets comprises receiving the
3 packets over the network from a plurality of different
4 sources, and wherein processing the header data comprises
5 selecting the group of packets for reassembly dependent
6 on the packets having been received from one or more
7 chosen sources among the plurality of different sources.

1 46. A method according to claim 45, wherein selecting
2 the group of packets comprises choosing the one or more
3 sources responsive to a level of reliability of a
4 connection over the network between the chose sources and
5 the network interface method.

1 47. A method according to claim 27, and comprising
2 writing the data packets that do not belong to the
3 identified group to the host memory substantially without
4 reassembly processing thereof by the network adapter.

1 48. A method according to claim 27, wherein writing the
2 fragments comprises writing the fragments by direct
3 memory access (DMA), substantially without involvement of
4 the host processor.

1 49. A method according to claim 27, wherein processing
2 the header data comprises determining a total length of
3 the data frame responsive to the header data of at least
4 one of the packets in the group, and wherein writing the
5 fragments comprises counting a quantity of the data in
6 the fragments reassembled in the data frame, and
7 comprising determining that all of the fragments have

41486S1

8 been reassembled by comparing the total length to the
9 quantity of the data reassembled.

1 50. A method according to claim 27, wherein processing
2 the header data comprises determining a total length of
3 the data frame responsive to the header data of at least
4 one of the packets in the group, and comprising writing a
5 frame header to the allocated space in the host memory,
6 indicating to the host processor a total length of the
7 data frame.